



NAVAL POSTGRADUATE SCHOOL

Monterey, California



THESIS

D2463

SIMPLIFICATION OF OBJECTS RENDERED
BY POLYGONAL APPROXIMATION

by

Lt. Michael Joseph DeHaemer, Jr., USN

December 1989

Thesis Advisor:

Michael J. Zyda

Approved for public release; distribution is unlimited

T249499

REPORT DOCUMENTATION PAGE

a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED		1b. RESTRICTIVE MARKINGS	
a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT	
b. DECLASSIFICATION/DOWNGRADING SCHEDULE			
d. PERFORMING ORGANIZATION REPORT NUMBER(S)		5. MONITORING ORGANIZATION REPORT NUMBER(S)	
5a. NAME OF PERFORMING ORGANIZATION Naval Postgraduate School	6b. OFFICE SYMBOL (If applicable) 37	7a. NAME OF MONITORING ORGANIZATION Naval Postgraduate School	
6c. ADDRESS (City, State, and ZIP Code) Monterey, CA 93943-5000		7b. ADDRESS (City, State, and ZIP Code) Monterey, CA 93943-5000	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER	
8c. ADDRESS (City, State, and ZIP Code)		10. SOURCE OF FUNDING NUMBERS	
		PROGRAM ELEMENT NO.	PROJECT NO.
		TASK NO.	WORK UNIT ACCESSION NO.
11. TITLE (Include Security Classification) SIMPLIFICATION OF OBJECTS RENDERED BY POLYGONAL APPROXIMATION (Unclassified)			
12. PERSONAL AUTHOR(S) DeHaemer Jr., Michael J.			
13a. TYPE OF REPORT Master's Thesis	13b. TIME COVERED FROM _____ TO _____	14. DATE OF REPORT (Year, Month, Day) 1989 December 21	15. PAGE COUNT 51
16. SUPPLEMENTARY NOTATION The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number) Graphics, Digitized Surfaces, Adaptive Subdivision, Data Reduction, Polygonal Approximation	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) Current Technology provides a means to obtain sampled data that digitally describes three-dimensional surfaces and objects. Three-dimensional digitizing cameras can be used to obtain sampled data that maps the surface of three dimensional figures and models. Data obtained from such sources enable accurate renderings of the original surface. However, the digitizing process often provides much more data than is needed to accurately recreate the surface or object. In order to use such data in real-time visual simulators, a significant reduction in the data needed to accurately render the sampled surfaces is required. The techniques presented were developed to drastically reduce the number of data points required to depict an object without sacrificing the detail and accuracy inherent in the digitizing process.			
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION	
22a. NAME OF RESPONSIBLE INDIVIDUAL Prof Michael J. Zyda		22b. TELEPHONE (Include Area Code) (408) 646-2305	22c. OFFICE SYMBOL 52ZK

Approved for public release; distribution is unlimited

**SIMPLIFICATION OF OBJECTS RENDERED BY
POLYGONAL APPROXIMATION**

by

Michael Joseph DeHaemer, Jr.
Lieutenant, United States Navy
B.S., University of Notre Dame, 1983

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN COMPUTER SCIENCE

from the

NAVAL POSTGRADUATE SCHOOL

December 1989

ABSTRACT

Current technology provides a means to obtain sampled data that digitally describes three-dimensional surfaces and objects. Three-dimensional digitizing cameras can be used to obtain sampled data that maps the surface of three dimensional figures and models. Data obtained from such sources enable accurate renderings of the original surface. However, the digitizing process often provides much more data than is needed to accurately recreate the surface or object. In order to use such data in real-time visual simulators, a significant reduction in the data needed to accurately render the sampled surfaces is required. The techniques presented were developed to drastically reduce the number of data points required to depict an object without sacrificing the detail and accuracy inherent in the digitizing process.

C.1 TABLE OF CONTENTS

I.	INTRODUCTION.....	1
A.	BACKGROUND	1
B.	THREE-DIMENSIONAL DIGITIZER.....	2
C.	SOLUTION TECHNIQUES	4
II.	ADAPTIVE SUBDIVISION	6
A.	BASIC METHOD.....	6
B.	SHORTCOMINGS OF THE BASIC METHOD	8
1.	Edge Gaps	8
2.	Simplistic Subdivision	9
3.	Aliasing	10
C.	VARIATIONS ON ADAPTIVE SUBDIVISION	11
1.	Filling the Edge Gap	12
2.	Number of Subpolygons.....	14
3.	Subdivision Locations.....	14
III.	POLYGON GROWTH	18
IV.	RESULTS	20
V.	FUTURE DIRECTIONS.....	35
A.	SEARCH.....	35
B.	SIMULATED ANNEALING	36
C.	HEURISTIC SEARCHES	38

VI. CONCLUSIONS.....	39
LIST OF REFERENCES.....	41
INITIAL DISTRIBUTION LIST	43

I. INTRODUCTION

As our technologically oriented civilization becomes increasingly more complex and sophisticated, the cost of training operators and technicians becomes more costly and time consuming. Inexpensive, three-dimensional simulators are important visualization tools that can provide an attractive alternative or supplement to traditional training methods [Ref. 1]. In order to be effective, these simulators must provide a sense of realism to the trainee; a real-time three-dimensional environment is an important contributor to this sense of realism. A need exists for realistic three-dimensional models of actual real world objects—ships, aircraft, automobiles and other vehicles—for use in these simulators. The use of a three-dimensional digitizing camera to obtain data from scale models can be used for this purpose. Such data provides an accurate depiction of objects used in the simulators [Ref. 2]. However, these images can contain several hundred thousand polygons—far too many polygons to be drawn in a real-time simulation without the use of very specialized and costly hardware. A semi-automated technique was sought for drastically reducing the number of polygons required to depict an object in real-time on relatively inexpensive hardware without sacrificing the detail and accuracy provided by the digitizing process [Ref. 3].

A. BACKGROUND

Many of the ongoing projects in the Graphics and Video Laboratory at the Naval Postgraduate School deal with the creation and use of low cost, real-time visual simulators [Ref. 1; Ref. 4; Ref. 5; Ref. 6]. Currently, simulators exist for land, air, sea and undersea vehicles which run on the Silicon Graphics IRIS-4D/GT series of

graphics workstations and make extensive use of the graphics capabilities of these machines. Recent work has focused on increasing the usefulness and effectiveness of the projects by incorporating real-world terrain data, including topography from digital survey databases, as well as reflectance and cultural feature information from photogrammetrically processed stereo-pair aerial photographs [Ref. 7]. These simulators, for example, now allow one to pilot a ship on the waters of the Sea of Japan, or to drive a jeep across the terrain of Fort Hunter-Liggett, California.

These efforts have been very important in adding to the visual realism of these simulators. However, there is still a need for realistic vehicles to inhabit them. Previously, vehicles were drawn up by hand on sheets of graph paper and painstakingly converted into polygonal representations for use in the programs. This technique is tedious, time consuming, error prone, and highly dependent upon the artistic talents of the individual creating the sketches. In order to overcome these problems, a project was undertaken to provide the vehicles in these simulators the same degree of accuracy and detailing as that available from using digital topography information for the terrain models.

B. THREE-DIMENSIONAL DIGITIZER

Cyberware Laboratories of Monterey, California produces a 3D digitizing camera that was used for data collection. It performs a cylindrical or linear scan of an object and produces a data file consisting of up to 1 million sample points that form a mesh that maps the surface of the object. In the case of the cylindrical scan, the mesh consists of radius values measured at uniformly spaced latitudes and longitudes around the object being digitized. The resulting mesh is identical to that produced by the lines of latitude and longitude drawn on a model globe. Borrowing from this analogy, it is convenient to visualize and discuss the data points in terms of their

latitude and longitude, and to use the terms north, south, east and west to express relationships among the sample points. The linear scans produce a mesh that resembles the grid on a sheet of graph paper; however, the terminology from the cylindrical scan is retained for consistency. It should be noted that the spacing between neighboring sample points is not necessarily the same along both latitudes and longitudes, nor are the number of samples the same in both directions.

A schematic diagram of the camera's operation is shown in Figure 1. A low power laser is directed through a series of lenses and mirrors to create a plane of light that is projected onto the subject. Mirrors view the reflected light from both sides of the subject in order to alleviate any problems caused by shadowing. These images of

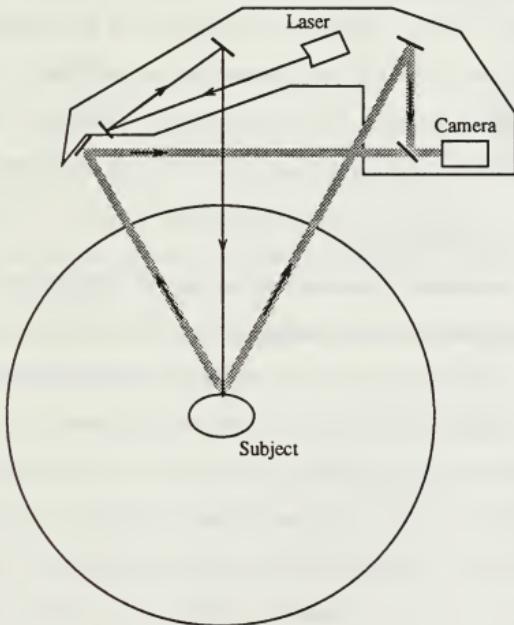


Figure 1. Schematic View of Cyberware Digitizer

the surface contour are combined at a half-silvered mirror and reflected into the lens of a high resolution CCD camera. The image is then processed by additional circuitry to extract the contour information and relay it to the controlling computer. The entire apparatus is mounted on a revolving framework so that it rotates around the subject creating a cylindrical scan of the surface contours. Alternately, small subjects can be moved past a stationary camera to produce linear scans of the subject. The most recent version of the camera is able to produce a 512×512 grid of sample points at a resolution of 0.7 millimeters in 15-20 seconds.

By using the shading and lighting models on the IRIS workstation, the renderings of these digitized objects are very realistic. We currently have scans of objects from simple geometric shapes to complex and detailed ornamental carvings, with an average of over 110,000 data points (see Photos 1-5 in Chapter IV). However the drawing capabilities of the IRIS workstations are not sufficient to allow models of this detail to be used in real-time 3D visual simulators. Therefore a method for reducing the number of data points required to accurately render an object was needed.

C. SOLUTION TECHNIQUES

There are two basic approaches to solving this problem, as shown in Figure 2. The first, and probably more straightforward, method is to systematically add or delete polygons from the object being rendered until the desired realism is obtained or the system's capabilities are exceeded. The second approach is to specify how many polygons the final object is to have, and then try to find the combination of polygons that gives the *best* rendering with that number of polygons. In either case, realism of the rendering must be balanced against drawing complexity and speed. The number of polygons to be drawn is a reasonable metric to be used in determining drawing complexity; however, a similar metric is not available to measure the realism of the

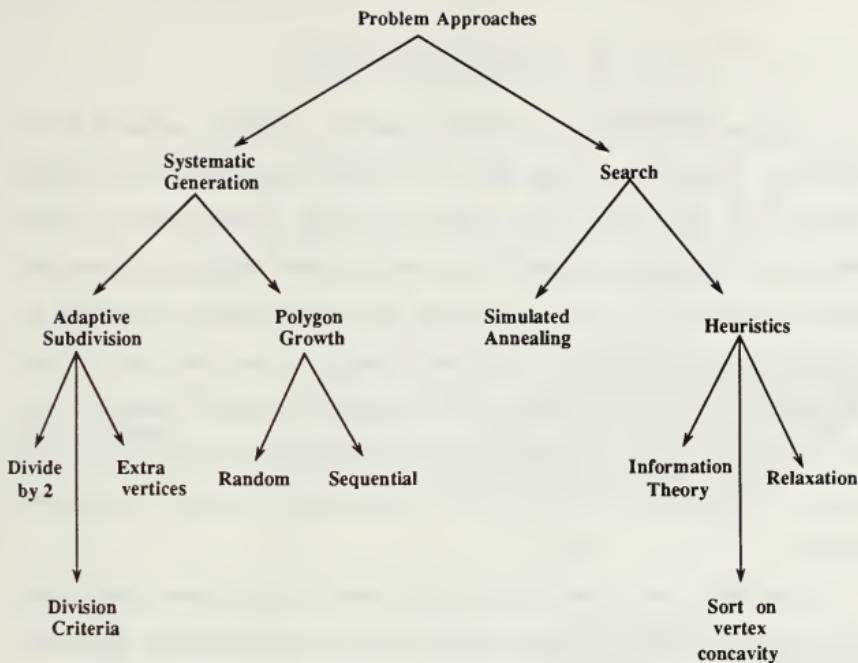


Figure 2. General Approaches to Data Reduction

rendering. For this reason, the techniques are not completely automated. The user is tasked with viewing the results of the data reduction and determining whether or not the realism is sufficient for the purpose at hand.

II. ADAPTIVE SUBDIVISION

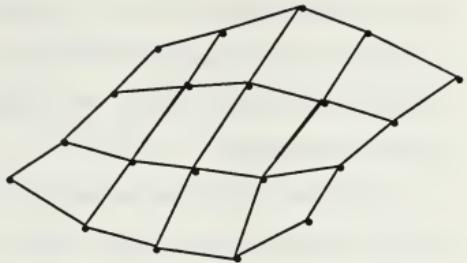
In their paper, Schmitt et. al. describe an adaptive subdivision method of fitting surfaces to sampled data [Ref. 8]. The technique approximates the surface represented by the sampled data points with bicubic Bernstein-Bézier surface patches. By constraining the patches to be continuous with neighboring patches, the coefficients of the patch can be determined. An accuracy metric is then used to measure the *closeness* of the approximating patch to the actual data. If the approximation is not within a user specified tolerance, the patch is subdivided into four smaller patches. The process is then recursively performed on these sub-patches until the set of patches approximates the sampled data to within the specified tolerance.

This technique has a useful property with respect to the simulator work being done in this laboratory—the method yields a reduction in the number of data points required to approximate the surface of a 3D object. By adjusting the method's tolerance, the user can make the approximation as accurate as necessary (limited by the precision of the digitizing process). However, the IRIS-4D/GT series of graphics workstations does not support the rendering of filled and shaded bicubic surfaces, instead the graphics engine supports high performance rendering of polygonal surfaces. For this reason, the methods presented here use simple polygons to approximate the surface of a 3D object.

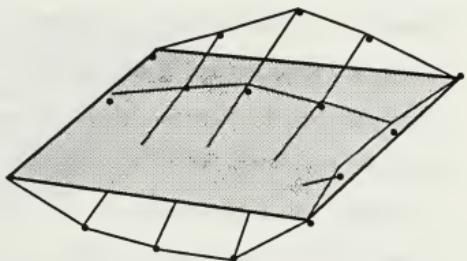
A. BASIC METHOD

The basic adaptive subdivision method as presented by Schmitt et. al. is sketched in Figure 3. Sketch 3a shows several data points produced by the sampling

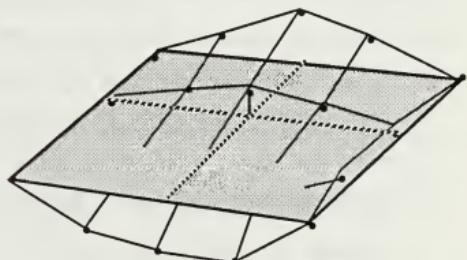
(a) Mesh of data points.



(b) Trial polygon.



(c) Measuring the errors.



(d) Trial polygon after subdivision

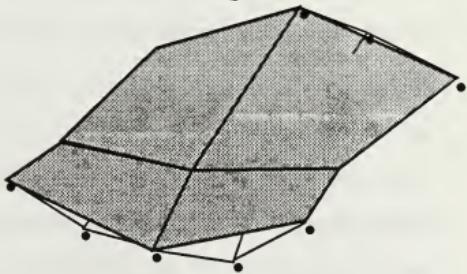


Figure 3. Surface-fitting Polygons with Adaptive Subdivision

process and sketch 3b adds a trial polygon that approximates the surface represented by the sample points. Due to the regularity of the mesh represented by the data points, the polygon is uniquely determined by noting the latitude and longitude of the northeast and southwest corners; these latitudes and longitudes form an index into the array of data points.

A bilinear interpolation is performed over the surface of the approximating polygon to locate the point on the surface 'directly below' a sample data point. The distance (or error) between the surface and the data point is then calculated, as shown in Figure 3c, and is compared to the user specified tolerance. If any of the errors exceeds the tolerance value then the polygon is divided into four smaller polygons, as shown in Figure 3d, and the process is repeated recursively.

When the approximation to the data is within the specified tolerance, the recursion is terminated and the polygon is saved in a linked list. Upon completion of the subdivision algorithm, the resulting polygons are displayed for the user to judge the results. Generally, the user runs the program repeatedly while varying the maximum tolerance until the results are suitable.

B. SHORTCOMINGS OF THE BASIC METHOD

Although the basic adaptive subdivision method provides a good starting point for this work, it has several shortcomings when polygons are used in place of bicubic patches in approximating the surface.

1. Edge Gaps

In [Ref. 8] the authors use bicubic Bernstein-Bézier surface patches to model the data. The transitions across patch edges to neighboring patches are guaranteed to be smooth and continuous due to the constraints applied when solving for the patch control points.

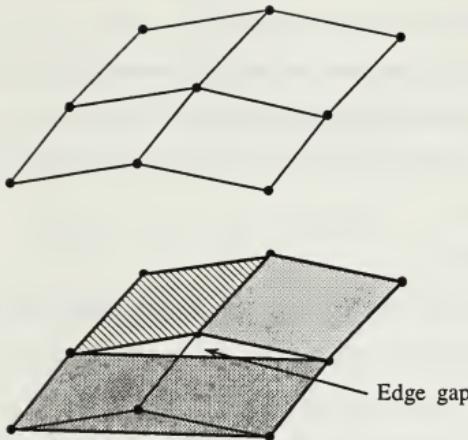


Figure 4. Edge Gaps

This is not the case when using polygons to approximate the data. As shown in Figure 4, it is very likely that the edges of neighboring polygons will not coincide. The result is that when the object is rendered using the polygonal approximation against a background of contrasting color, holes can appear where these edge gaps exist.

2. Simplistic Subdivision

The algorithm presented by Schmitt, et. al. always subdivides the surface patch into four smaller patches along the latitude and longitude passing closest to the center of the patch. Such subdivision proved sufficient for their purposes and results in a simple and fast procedure. However, when used for surface-fitting with polygons, it can leave artifacts in the final rendering. Additionally, the characteristics of the sampled data are not used to intelligently guide or optimize the subdivision process.

The artifacts left by the basic subdivision method are similar to the creasing problem encountered with some methods for creating fractal mountains [Ref. 9 and Ref. 10]. Due to the regularity of choosing the division point of the polygons, the edges and corners of adjacent polygons tend to line up with each other. This can create 'creases' in the surface of the final rendered object, causing the surface to appear unnatural.

3. Aliasing

There are two forms of aliasing that can occur when rendering the reduced data. The first results when an edge gap is viewed tangentially (Photo 1). This causes the outline of the rendered object to have a stair-step appearance. Data sets generated with large accuracy tolerances are much more susceptible to this effect than those with small tolerances.

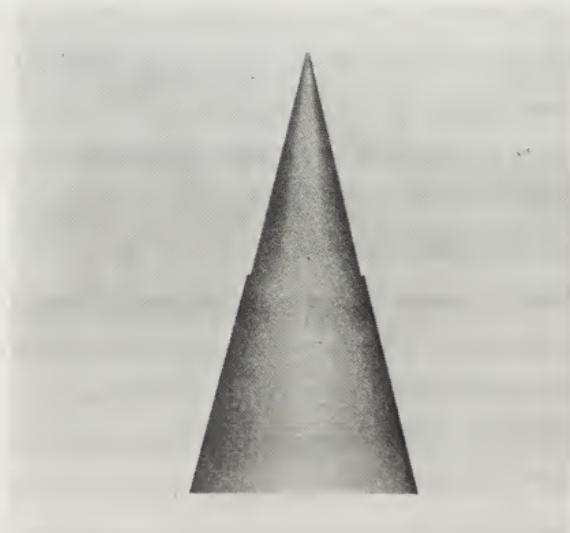


Photo 1. Aliasing Due to Edge Gaps.

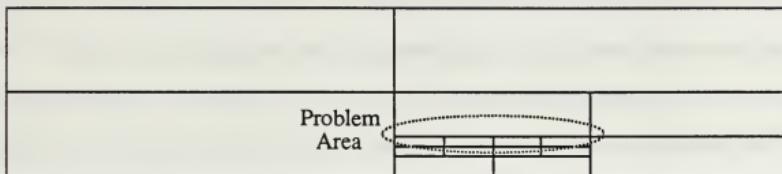


Figure 6. The cause of aliasing

The second form of aliasing results from the use of the IRIS's lighting models. The raw data produced by the Cyberware digitizer represents a mesh of sampled data points in a 512×512 grid. However, the volume scanned by the digitizer is usually taller than the subject. This causes the highest and lowest latitudes of the data grid to contain zero or meaningless values. Such latitudes are usually trimmed off when the raw data is edited and the data mesh is no longer square. In such a case, many of the approximating polygons become very long and narrow, as shown in Figure 6. The shading routines essentially calculate the correct shade at the vertices of the polygon and then use interpolation to determine the shading of the interior of the polygon. When several of these strip-like polygons are adjacent to a larger polygon, there can be a noticeable step in shading between the small polygons and the larger polygon. This problem can be minimized by using material definitions with low values of specular reflectance [Ref. 11] and by using small tolerances.

C. VARIATIONS ON ADAPTIVE SUBDIVISION

As described above, there are several weaknesses in using the basic adaptive subdivision method for building polygonal approximations to sampled surface data. In

order to overcome these weaknesses and to investigate alternative strategies, several variations were developed. Each variation is coded as a separate module, and the modules can be combined in many ways as shown in Figure 7— one module from each level. There are three basic ways in which the subdivision procedure can be modified—keep extra vertices or extra polygons to fill the edge gap, divide a polygon into two parts instead of into four parts, or change the way in which a polygon is actually subdivided.

1. Filling the Edge Gap

The first level of variations deals with the problem of edge gaps appearing in the reduced-data renderings. These gaps can be dealt with in several different ways. First, depending upon the accuracy required of the final rendering and the texture of the subject's surface, the edge gaps may not be visible and therefore no corrective action is necessary. Thus the *corners_only* module makes no attempt to fill the edge gaps, and was used in the original program that pointed out the edge gap problem.

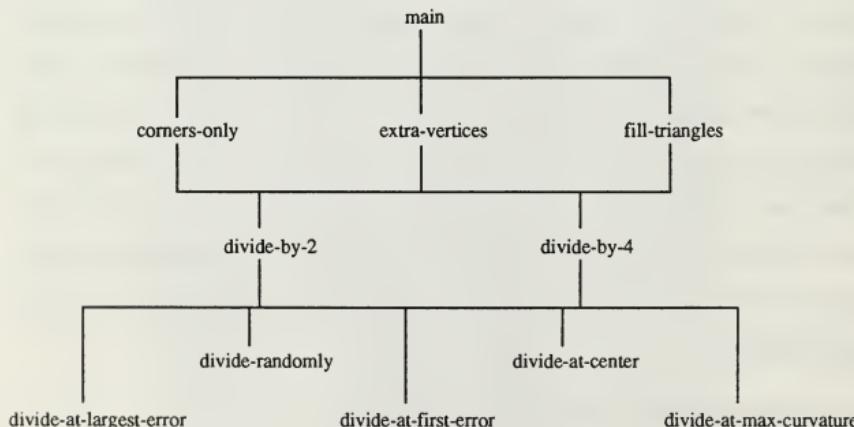


Figure 7. Adaptive Subdivision Variation Hierarchy

However, most of the test cases have demonstrated the need for filling these edge gaps.

One possible solution is to create additional polygons to fill these gaps as illustrated in Figure 8a. These extra polygons are created by connecting vertices lying along a common latitude or longitude with vertices that are also the corners of an adjacent polygon. The disadvantages of this solution are that more polygons are created which raises drawing complexity and these fill polygons lie in a plane perpendicular to the object's surface.

The alternative solution to the problem of edge gaps is to retain extra vertices along the edge of a polygon as needed to keep the edge gaps from appearing. In Figure 8b, the larger polygon contains the center vertex in the mesh as well as its four corner points. Note that only those vertices along an edge that are also corner vertices of neighboring polygons need to be retained. This approach has the drawback

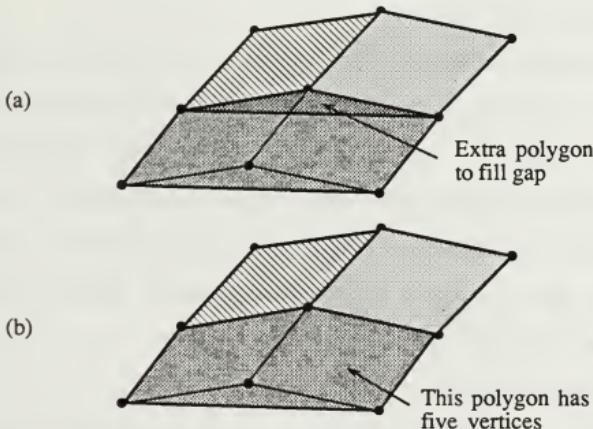


Figure 8. Filling the Edge Gap

of increasing the average number of vertices per polygon in the rendered object, but also has several benefits. First, little extra data storage is required since the extra vertices in the edge of one polygon are the corner vertices of a neighboring polygon. Second, the lighting capabilities of the IRIS automatically fill the gaps with the correct shade depending on material, light(s), etc. And third, because vertex normals are used with the lighting model, the shading appears smooth across polygon boundaries.

2. Number of Sub-Polygons

In the original version of the polygon-surface-fitting procedure, a polygon was always subdivided into four smaller parts. This often resulted in long thin polygons that degraded the appearance of the rendered image. As a means to alleviate this problem, a variation of the algorithm was programmed to divide a polygon into two sub-polygons. In this case, the decision to divide along a latitude or a longitude could be based on which dimension of the polygon is greater—north-south or east-west.

3. Subdivision Locations

The remaining variations of the adaptive subdivision technique all concern the choice of how to subdivide a polygon that does not meet the accuracy specified by the user. This accuracy is measured as the distance from the surface of the approximating polygon to each sample data point being approximated. In each case, the accuracy of all the data points 'over' the polygon are checked; as soon as one distance exceeds the tolerance, the subdivision method is invoked to determine where the polygon should be split.

a. *Center Subdivision*

This is the subdivision method used by the original adaptive subdivision technique. No effort is made to analyze the data points in order to intelligently choose

a point through which to subdivide a polygon. Instead the algorithm simply finds the center latitude and longitude of the polygon and divides it along those lines.

b. First Error

In determining the accuracy with which a polygon represents a group of sampled data points, the distance from the polygon's surface to each data point is measured and compared to the tolerance. Depending upon the dimensions of the polygon, these points are checked in a north-to-south/west-to-east or a west-to-east/north-to-south manner. The checking of the data points terminates when a measurement exceeds the tolerance, and the latitude and longitude of the offending point are used to divide the polygon.

c. Greatest Error

This version is similar to *first_error*, except that accuracy measurement does not end when an error measurement exceeds the user specification. Instead, the process is continued so that the point of maximum error can be found. Upon completing the accuracy check, the latitude and longitude of the point of maximum error are used to subdivide the original polygon. Note that this slows down the procedure due to the large number of comparisons which must be made.

d. Maximum Curvature

In attempting to intelligently choose a point to divide a polygon, consider how it might be done manually. An obvious place to put an edge of a polygon is along an apparent edge on the surface—such as along the edge of a cube. For a computer algorithm, detecting these edges is not a simple task. However, for the adaptive subdivision to work, the exact edges are not needed.

Consider the mesh of data points resulting from digitizing the surface of a cube and picture the normal vectors associated with each such data point. The data points that lie on the top side of the cube will have normal vectors that are all parallel

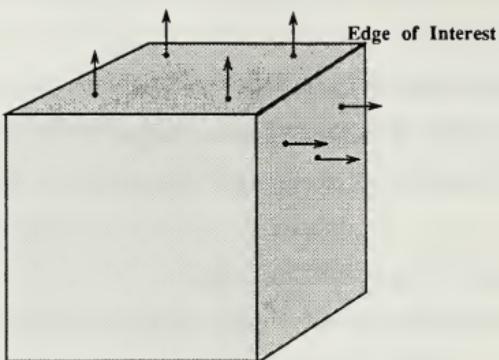


Figure 9. Normal Vectors on Opposite sides of an Edge

to each other (or nearly parallel). As shown in Figure 9, this will not be the case for the normal vectors of data points that lie on opposite sides of an edge. Therefore, the angle formed between the normals of nearby data points gives an indication of how sharply the surface is curving.

In this method, before the subdivision begins, each data point is assigned two 'curvature' measures—one in the north-south direction and the other east-west. These measures are based on the length of the vector resulting from the subtraction of the normal vectors for the data points on either side of the data point being considered (see Figure 10). Then during the subdivision process, the location of maximum 'curvature' is found in much the same way as the point of maximum error was located above. The latitude and longitude of this point then become the dividing lines for the new polygon. Because one of the points at the corner of a polygon will always have the highest curvature (since it was chosen as the place to subdivide the

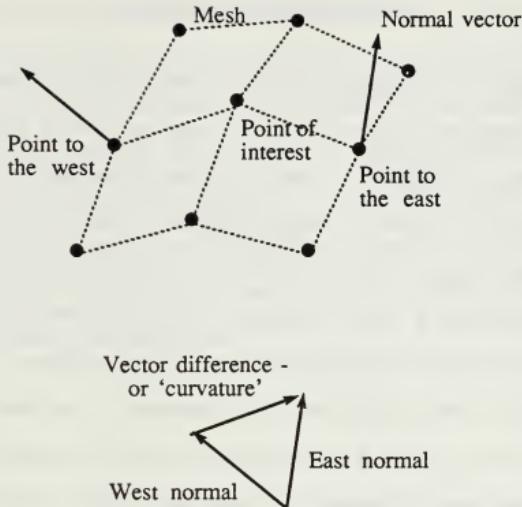


Figure 10. Vector Difference to represent Curvature

previous polygon), the corners should be excluded from the scan in order to prevent an infinite recursion.

III. POLYGON GROWTH

Another method to systematically generate surface-fitting polygons is to use a *polygon growth* technique (Figure 2). First a *seed* polygon is selected from the original set of data, as in Figure 11a. Then a neighboring polygon is selected and the two are *combined* into a larger trial polygon (Figures 11b and c). If the trial polygon passes the accuracy metric, additional neighbors are selected in an attempt at further growth, as in 11d through g. When all attempts at further growth fail, the polygon is added to the list of polygons describing the reduced object, as shown in Figure 11h. A new *seed* polygon is selected and the process repeated. When there are no more polygons from which to choose a new seed, the process is terminated. Variations on this method include using a random or a sequential selection strategy for the next *seed* polygon.

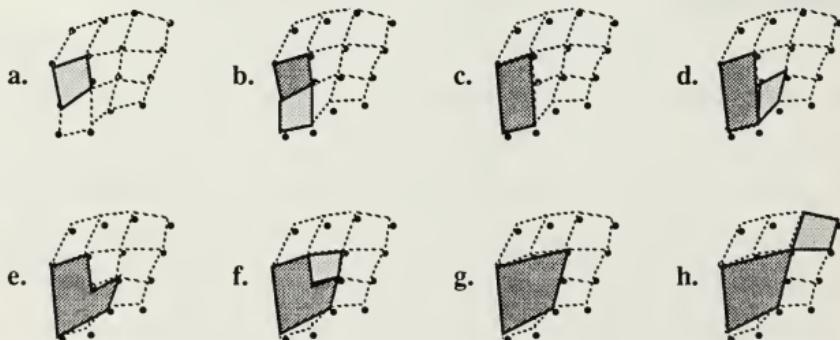


Figure 11. Polygon Growth

This method of fitting polygons to the sampled data did not prove as useful as the adaptive subdivision methods, for several reasons. First, it is an essentially brute force method of surface fitting and very repetitive in nature. Many trial polygons are created, only to be 'thrown out' because they exceed the error tolerance. This causes the procedure to execute quite slowly—orders of magnitude more slowly than the adaptive subdivision methods. Second, the data is not significantly reduced. Many of the generated polygons are of complicated shapes and use a large number of vertices. Thus many of the original sample data points are retained, increasing storage requirements and slowing drawing times.

IV. RESULTS

This project was undertaken to investigate possible methods for creating realistic ships, trucks and other inhabitants of the real-time 3D visual simulators being developed at the Naval Postgraduate School. The majority of the research efforts were directed at developing the adaptive subdivision techniques, since this area seemed most promising for near term results.

The general test procedure consisted of viewing renderings of both the full data set and of a reduced data set. The data reduction program was repeated while the tolerance was adjusted based on trial-and-error and experience in order to achieve a good balance between the accuracy of the rendering and its drawing complexity. This process was then repeated for each of the variations of the adaptive subdivision method. The quality of the final renderings is necessarily a subjective judgement made by the user; however, every effort was made to be consistent across all of the reduction variations for each test subject.

Tables 1-5 give the results from applying the adaptive subdivision variations discussed in Chapter II to a number of different digitized test objects. Table 1 contains the results from a digitized bust of Spock. Tables 2 and 3 are from a motorcycle helmet and a paint tube respectively; while the last two tables are from linear scans of a telephone and of an ornamental carving. Photo 2a shows a rendering of Spock's bust with the full set of sampled data points while Photo 2b shows the rendering of the reduced data. Similarly Photos 3-6 show before and after renderings of the other test objects. Lastly Photos 7 and 8 show the wireframe of the reduced paint tube and carving respectively.

Each table records various measures of the algorithm's performance. The uppermost row lists values for the raw data, and each of the other rows gives the results after applying the variation indicated. *Spec* is the tolerance value used with that particular adaptive subdivision method. The columns labeled *polys* and *vrtcs* list the number of polygons and the number of vertices retained from the original sampled data set. Since in the *extra vertex* and *fill polygon* variations of the adaptive subdivision technique, the polygons may not have four vertices per polygon, the number of polygons and the number of vertices only give a partial evaluation of the complexity of the final rendered object. In order to get a better estimation of complexity, the number of calls to the IRIS's graphics library are also tabulated (*gl calls*).

The data structures used for storing the object consists of a table of vertices and a linked list of polygons. Each vertex structure in the table contains fields for the x/y/z-coordinates and the i/j/k-components of the unit normal vector at that vertex. The polygon structures are maintained in a linked list and contain a pointer to an array of pointers to entries in the vertex table. The *mem* column in the tables give the number of bytes needed to store such a data structure.

The *fill polygon* variation builds small polygons to fill in the edge gaps in between the other polygons. The *f_p* and *f_v* fields of the tables list the number of extra polygons used, and the number of extra pointers into the vertex table that are required. Thus they give an indication of the extra graphics library function calls required of this variation. Note that the *f_p* polygons are also included in the *poly* total.

The adaptive subdivision variations themselves are listed down the left side of the charts. The first letter of each code designates how the subdivision point is chosen—*c* is center; *f* is first error; *g* is for greatest error; *r* is randomly selected, and

m is for point of maximum curvature. The remaining letters indicate the other levels in the variation hierarchy (Figure 2)—*e* designates that extra vertices are retained, while *f* and 2 indicate creating fill polygons and subdivision-by-two respectively. There are no entries for the combination of divide at point of first error and subdivision by two. This is because the location where the error measure exceeded the specified tolerance was very often along one edge of the polygon; effectively subdividing by two. The renderings of the elementary variations—*c*, *f*, *g*, *r* and *m* as well as their divide-by-two variations are generally of poor quality due to the presence of gaps. However, the data is included as a reference for measuring the overhead caused by the edge gap treatments.

The general results indicate that for data from objects with rough, complex surfaces, the best results are obtained when the strategy of subdivision at the point of greatest error is used, and extra vertices are retained in order to fill any edge gaps—*ce* or *c2e*. This is not surprising, since such subdivision tends to reduce the measured error quickly. However, in some cases, this method results in a loss of detail in areas where the surface has sharp edges. This is readily apparent near the button edges in the linear scan of the telephone (Photo 5). In these cases, the center division variations *ce* and *c2e* can sometimes give better results, as in the case of the ornamental carving (Table 5 and Photo 6).

Better results were expected from the strategy of dividing at the point of maximum curvature—*me*. The idea seems to work well in areas of rapidly curving surfaces, but begins to suffer in areas that are relatively smooth and flat. In such areas, the difference between adjacent normal vectors tends toward very small values, possibly smaller than the digitizing camera's resolution, causing the choice of the subdivision point to become more random. A hybrid technique that uses curvature

based subdivision for the first few levels of recursion and then switches to a greatest error technique when the surfaces become nearly flat may improve these results.

spec	polys	vrtcs	gl calls	mem	f_p	f_v	notes
all	—	112128	112640	1121280	6293504		
c	0.047	14004	16604	140040	857024		h
ce	"	"	"	149226	875396		
cf	"	14872	"	147190	878284	868	2711
c2	0.035	16043	19747	160430	1002120		h a
c2e	"	"	"	173792	1028844		a
c2f	"	17321	"	170872	1033228	1278	3943
f	0.080	23595	33770	235950	1606220		h
fe	"	"	"	275064	1684448		
ff	"	32065	"	304248	1810576	8470	25679
g	0.026	15970	21827	159700	1058316		h
ge	"	"	"	181998	1102912		
gf	"	18036	"	176980	1109404	2066	6574
g2	0.030	12821	17918	128210	660692		h
g2e	"	"	"	147476	899224		
g2f	"	14647	"	143344	905568	1826	5741
r	0.025	30069	42324	300690	2027004		h a
re	"	"	"	348128	2121880		a
rf	"	33160	"	327042	2104436	3091	10085
r2	0.029	21124	30339	211240	1440964		h
r2e	"	"	"	246532	1511548		
r2f	"	24972	"	243544	1536356	3848	12304

first letter: divide at center, first error, greatest error,
maximum curvature or random point

other letters: fill edge gaps with extra vertices or fill polygons,
and divide-by-2 instead of by four

notes: holes due to edge gaps
artifacts of long narrow polygons

Table 1. Results of Data Reduction on Bust of Spock

(a) Full data set



(b) Reduced data set



Photo 2. Bust of Spock in Profile

spec	polys	vrtcs	gl calls	mem	f_p	f_v	notes
all		93184	93696	931840	5232640		
c	0.055	4730	5664	47300	291032		h
ce	"	"	"	50474	297380		
cf	"	4956	"	49156	296552	226	706
c2	0.045	6544	8191	65440	412580		h
c2e	"	"	"	71236	424172		
c2f	"	6841	"	67846	419768	297	906
f	0.043	11200	16105	12800	766780		h
fe	"	"	"	131080	803340		
ff	"	15280	"	144864	862908	4010	12032
g	0.040	4362	6238	43620	296800		h
ge	"	"	"	50648	310856		
gf	"	4897	"	48054	309948	535	1682
g2	0.035	4544	6586	45440	311640		h
g2e	"	"	"	53074	326908		
g2f	"	5190	"	50740	327408	646	2004
r	0.070	6598	10419	65980	476476		h
re	"	"	"	80810	506136		
rf	0.040	11783	16127	115976	777772	1262	4121
r2	0.035	8933	13124	89330	617596		h
r2e	"	"	"	105258	649452		
r2f	"	10510	"	102606	656764	1577	5061

first letter: divide at center, first error, greatest error,
maximum curvature or random point

other letters: fill edge gaps with extra vertices or fill polygons,
and divide-by-2 instead of by four

notes: holes due to edge gaps
artifacts of long narrow polygons

Table 2. Results of Data Reduction Motorcycle Helmet

(a) Full data set



(b) Reduced data set



Photo 3. Motorcycle Helmet

spec	polys	vrtcs	gl calls	mem	f_p	f_v	notes
all	75264	75776	752640	4229120			
c	0.058	3233	4017	32330	203000		h
ce	"	"	"	35070	208480		
cf	"	3425	"	33940	207756	192	613
c2	0.056	4621	5916	46210	295036		h
c2e	"	"	"	50800	304216		
c2f	"	4849	"	48052	300544	228	693
f	0.061	4136	5871	41360	280196		h
fe	"	"	"	47628	292732		
ff	"	5116	"	49420	304156	980	3050
g	0.035	2963	4366	29630	205212		h
ge	"	"	"	34966	215884		
gf	"	3633	"	33028	215208	400	1299
g2	0.030	3617	5407	36170	252672		h
g2e	"	"	"	43006	266344		
g2f	"	4170	"	40720	266212	553	1726
r	0.030	9995	14841	99950	695408		h
re	0.042	8176	12367	98056	607796		
rf	"	9597	"	93794	610640	1421	4596
r2	0.045	7400	11632	74000	532896		h
r2e	"	"	"	90208	565312		
r2f	"	9319	"	90188	580624	1919	6175

first letter: divide at center, first error, greatest error,
maximum curvature or random point

other letters: fill edge gaps with extra vertices or fill polygons,
and divide-by-2 instead of by four

notes: holes due to edge gaps
artifacts of long narrow polygons

Table 3. Results of Data Reduction on Paint Tube

(a) Full data set



(b) Reduced data set



Photo 4. Tube of Paint

spec	polys	vrtcs	gl calls	mem	f_p	f_v	notes
all		73920	74360	739200	4151840		
c	0.010	19631	23028	196310	1194452		h
ce	"	"	"	208632	1219096		
cf	"	20396	"	202870	1213692	765	2515
c2	0.010	18520	22611	185200	1151668		h
c2e	"	"	"	200188	1181644		
c2f	"	19746	"	195416	1181908	1226	3882
f	0.015	11884	15809	118840	775404		h
fe	"	"	"	133774	805272		
ff	"	13611	"	133622	818864	1737	5654
g	0.010	11801	15897	118010	775544		h
ge	"	"	"	133800	807124		
gf	"	13046	"	128630	806744	1245	4065
g2	0.010	10268	14046	102680	680792		h
g2e	"	"	"	117130	709692		
g2f	"	11529	"	113406	712332	1261	4102
r	0.011	18788	25875	187880	1250564		h
re	"	"	"	215326	1305456		
rf	"	20459	"	202190	1292552	1671	5484
r2	0.015	14249	20460	142490	971852		h
r2e	"	"	"	173936	1063344		
r2f	"	17480	"	170858	1077340	2519	8105

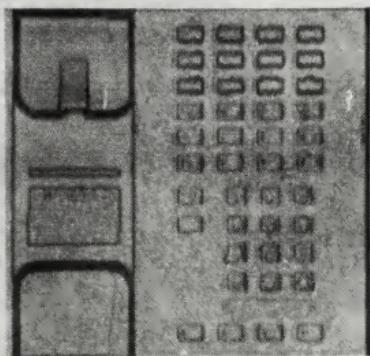
first letter: divide at center, first error, greatest error,
maximum curvature or random point

other letters: fill edge gaps with extra vertices or fill polygons,
and divide-by-2 instead of by four

notes: holes due to edge gaps
artifacts of long narrow polygons

Table 4. Results of Data Reduction on Linear Scan of Telephone

(a) Full data set



(b) Reduced data set

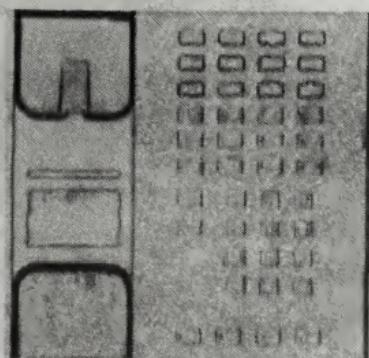


Photo 5. Telephone

spec	polys	vrtcs	gl calls	mem	f_p	f_v	notes
all		130560	131022	1305600	7325696		
c	0.032	15946	19447	173078	1018240		h
ce	"	"	"	159460	991004		
cf	"	17013	"	168332	1017284	1067	3369
c2	0.032	16818	21047	168180	1060220		h
c2e	"	"	"	184700	1093260		
c2f	"	18181	"	179372	1093508	1363	4233
f							u
fe							u
ff							u
g	0.032	11875	17009	118750	808752		h a
ge	"	"	"	138936	849124		a
gf	"	13098	"	129080	839196	*	*
g2	0.032	11278	15065	112780	761804		h a
g2e	"	"	"	131088	799420		a
g2f	"	13008	"	127386	805856	1730	5573
r	0.031	21726	31302	217260	1484784		h
re	"	"	"	255182	1560628		
rf	"	24520	"	241122	1534860	2794	9137
r2	0.031	18249	26805	182490	1261512		h
r2e	"	"	"	216324	1329180		
r2f	"	21820	"	212768	1350636	3571	11568

first letter: divide at center, first error, greatest error,
maximum curvature or random point

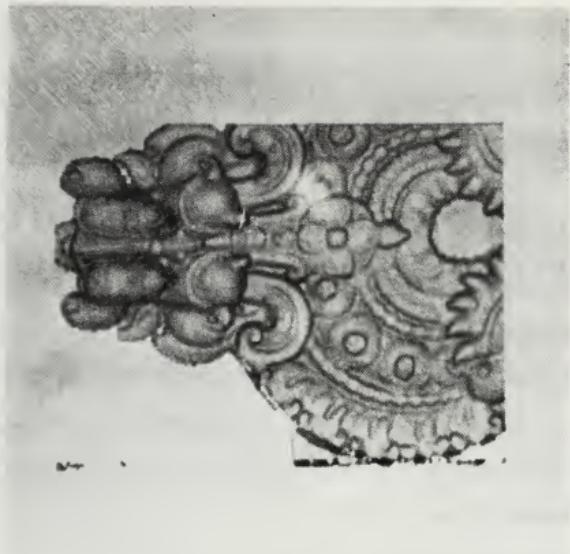
other letters: fill edge gaps with extra vertices or fill polygons,
and divide-by-2 instead of by four

notes: holes due to edge gaps
artifacts of long narrow polygons

unable to get adequate rendering

Table 5. Results of Data Reduction on Linear Scan of Ornamental Carving

(a) Full data set



(b) Reduced data set

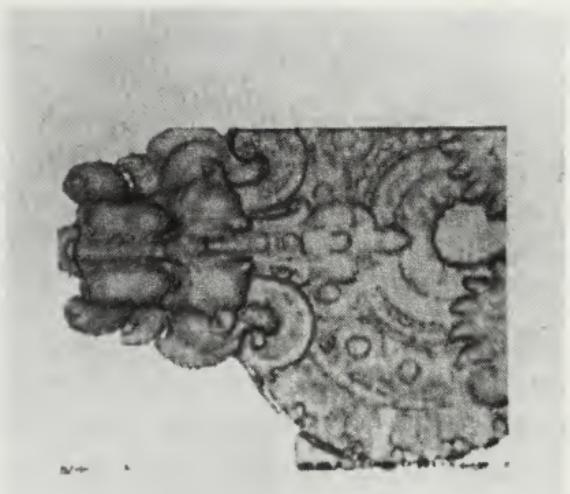


Photo 6. Ornamental Carving

(a) Wireframe of
paint tube



(b) Wireframe of
ornamental carving

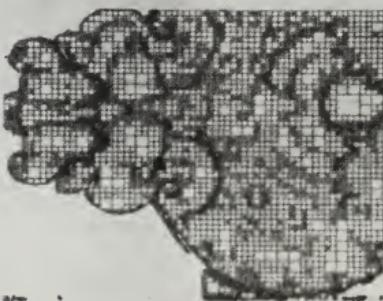


Photo 7. Example Wireframes

V. FUTURE DIRECTIONS

The polygon-fitting techniques discussed thus far create a polygonal surface that approximates sampled data by ‘adding’ polygons until the approximation falls within a user specified tolerance. The user can control the accuracy of the final rendered object, but cannot directly control the final number of polygons (by running the procedure repeatedly and ‘tweaking’ the tolerance, the user can zero in on the desired number of polygons). In many cases, it would be desirable to specify that the final approximation consist of a given number of polygons and have the data reduction procedure return the *best* approximation using that number of polygons.

This second approach to the problem can be viewed as a search. The nodes of the search correspond to sets of P polygons to be used in the rendering, and the search space consists of all possible combinations of those P polygons out of the N original polygons. The goal of the search then is to find the set of polygons that *best* approximates the original data. Because the problem is combinatorial in nature (i.e. the search space grows as the factorial of N) brute force search techniques would not be feasible for most real world problems, therefore techniques borrowed from the artificial intelligence field such as simulated annealing and the use of heuristics are to be investigated [Ref. 12 and Ref. 13].

A. SEARCH

In trying to find the *best* approximation to the original data, one is trying to minimize an error function (i.e. maximize the accuracy). One approach to minimizing a function is to choose a possible solution and evaluate the function. By carefully changing the function’s arguments in small increments, its value can be made to

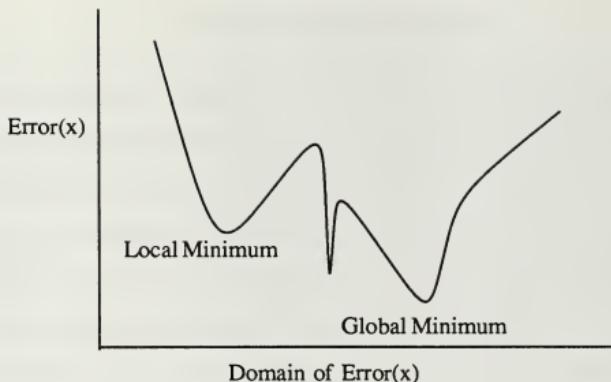


Figure 12. Local Minima in an Error Function

decrease. When no further changes to the arguments results in a decrease in the function's value, a minima has been found.

A problem that arises with this incremental improvement strategy is that this procedure can become trapped at a local minimum, and therefore never find the global minimum (Figure 12). This is because since only changes that move closer to the goal (minimum error) are allowed, there is no way of climbing over a small 'hump' in the function's value. One way to minimize this effect is to repeat the incremental improvement process many times, each from a different starting points and simply retaining the best solution. However, this can fail if the minimum is located in a narrow 'valley' in the error curve.

B. SIMULATED ANNEALING

Simulated annealing is similar to the incremental improvement technique except that 'uphill' excursions are allowed in a controlled way. An initial guess for the best

solution is formed and evaluated. The best solution is then randomly permuted in an effort to improve it. If the permutation results in a lower error, then it becomes the new best solution and the permutation/evaluation cycle is repeated. If the new guess has a higher error, a probability function is used to determine whether to replace the current best solution.

The probability is determined by a system parameter called the system *temperature*. During the initial stage of the algorithm, the temperature is high and the probability that a more costly solution is allowed to replace the current solution is also high. As the process continues, the temperature is gradually reduced and the system ‘cools’. This results in a lower probability of keeping a solution that is ‘uphill’ from the current one. Eventually, the system *freezes* and no more changes are possible; the current guess is then the solution.

At high temperatures, the solution is allowed to make many long random uphill jumps, effectively conducting a broad survey of the search space. As the system cools, the uphill jumps get progressively smaller and less frequent until the system freezes. Since all better solutions are accepted and only some of the poorer guesses are accepted, the overall trend is a movement toward smaller errors. The occasional ‘uphill’ excursion helps keep the solution from getting stuck at a local minimum.

Simulated annealing does not guarantee the most optimal solution, but it has been used effectively to obtain near-optimal solutions for very large and complex problems [Ref. 12]. By choosing an initial set of polygons with which to render a digitized object and permuting this set in a consistent way, this technique may be useful in data reduction efforts.

C. HEURISTIC SEARCHES

Stochastic search techniques such as simulated annealing examine a large number of states in the search space, and because each candidate successor state is randomly generated, they spend a great deal of time examining candidates that are worse than the current best ‘solution’. In fact, at low system temperatures, simulated annealing programs spend a large fraction of their cycles examining these worse states. A heuristic search is designed to minimize examining extra states and instead to concentrate on search paths that look particularly promising. The next candidate solution to process is picked based on some rules that weigh the costs and benefits of the different possible choices. At each stage, the best looking candidate is chosen. At the same time, it is possible to eliminate or *prune* candidates that look bad, or look worse than the current solution by some threshold value.

If we consider the solution to be a set of P polygons, then our method for generating candidate solutions might be to exchange a polygon in our set P with one not in the set. A possible heuristic to use in choosing which polygons to exchange would be to swap out the polygon with the largest curvature (i.e. is most nearly flat) with one that has a small curvature. Similarly, we may instead chose to rank the edges or vertices to discover which vertices, edges, and polygons carry more meaning, etc. The idea is to use one’s knowledge about the system and the problem in order to more quickly guide the program to the solution.

VI. CONCLUSIONS

Current work on this project has focused on the systematic generation approaches to the simplification of rendered objects. Both adaptive subdivision and polygonal growth methods have been implemented, along with several of the variations. The initial results are very encouraging; sample data of a human bust consisting of over 110,000 points was successfully reduced to just over 9,000 polygons using the extra vertex-greatest error variation of the adaptive subdivision method (ge) (as in Figure 8). This drastic reduction in object complexity resulted in an order of magnitude increase in drawing speeds. Further investigation into improving the accuracy metrics, and into rules concerning how to subdivide a polygon should result in greater reductions.

Other issues include:

- loss of features that are finer than the digitizing camera's resolution,
- methods of combining several sets of data from the camera for objects too large to digitize in one pass,
- being able to have different accuracy metrics for different parts of the same object (e.g. when working with a digitized bust, may want to require higher accuracy in the area of the face than in the back of the head),
- automatically adjusting the tolerance based on surface features,
- using surface color information that will be provided by the next generation of the Cyberware Laboratories rapid digitizer,
- defining boundaries for areas of an objects surface, so that simplification occurs within these boundaries. This would allow assignment of colors to parts of the object without fear of a polygon crossing from one color to another,
- eliminating the artifacts produced by the subdivision methods,
- improving execution times of the polygonal growth technique,
- investigate various data structures for storing and manipulating the objects.

These issues represent areas of current investigation; for the future, research into the use of search techniques will be expanded. Small scale programs written to learn and experiment with these technique need to be scaled up and adapted to the data reduction problem. Effective heuristic rules need to be devised in order to allow such techniques to be used efficiently, and ways of measuring how accurate a rendering will look to a user need to be investigated.

Our goal is to be able to use these processes to create models for the simulators developed here at NPS. By using the digitizing process on scale models, the design, creation, and coding of vehicles and other objects for these simulators can be made easier, faster, and more realistic. Initial results are good, indicating that this is a promising area of investigation.

REFERENCES

- 1 Zyda M., McGhee R., Ross R., Smith D., and Streyle D., "Flight Simulators for Under \$100,000," *IEEE Computer Graphics and Applications*, Vol. 22 No. 1, pp. 19-27, January 1988.
- 2 Addleman D. and Addleman L., "Rapid 3D Digitizing," *Computer Graphics World*, Vol. 8 No. 11, pp. 41-44, November 1985.
- 3 Akeley K. and Jermoluk T., "High Performance Polygon Rendering," *Computer Graphics*, Vol. 22 No. 4, pp. 239-246, August 1988.
- 4 Weeks, G. K. and Phillips, C. E., *The Command and Control Workstation of the Future—Subsurface and Periscope Views*, Masters Thesis, Naval Postgraduate School, Monterey, California, June 1989.
- 5 Naval Postgraduate School Technical Report NPS52-88-022, *Graphics Workstations and 3D Visual Simulation: Some Performance Expectations and Measurements*, by Zyda, M. J., Fichten, M. A. and Jennings, D. H., July 1989.
- 6 Strong, R. P., Winn, M. C., *The Moving Platform Simulator II: A Networked, Real-Time Visual Simulator with Distributed Processing and Line-of-Sight Display*, Masters Thesis, Naval Postgraduate School, Monterey, California, June 1989.
- 7 Breden, W. O. and Zanoli, J. J., *Visualization of High Resolution Digital Terrain*, Masters Thesis, Naval Postgraduate School, Monterey, California, June 1989.
- 8 Schmitt F. J. M., Barsky B. A., and Wen-Hui Du, "An Adaptive Subdivision Method for Surface-Fitting from Sampled Data," *Computer Graphics*, Vol. 20 No. 4, pp. 179-188, August 1986.
- 9 Barnsley, M. F. and others, "The Science of Fractal Images," Springer-Verlag, 1988.
- 10 Miller, G. S. P., "The Definition and Rendering of Terrain Maps," *Computer Graphics*, Vol. 20 No. 4, pp. 39-48, August 1986.
- 10 Kirkpatrick S., Gelatt C. D., Jr., and Vecchi M. P., "Optimization by Simulated Annealing," *Science*, 220:671-680, 1983. Reprinted in *Neurocomputing*, pp.554-568, MIT Press, 1988.

- 11 Silicon Graphics, Inc., *GT Graphics Library User's Guide*, 1988.
- 12 Korf R. E., "Search: A Survey of Recent Results," *Exploring Artificial Intelligence*, pp.197-239, H. E. Shrobe editor, Morgan Kaufmann Publishers Inc. 1988.
- 13 Farin, G., *Curves and Surfaces for Computer Aided Geometric Design—A practical Guide*, Academic Press, Inc., 1988.

INITIAL DISTRIBUTION LIST

- | | | |
|----|--|----|
| 1. | Defense Technical Information Center
Cameron Station
Alexadria, VA 22304-6145 | 2 |
| 2. | Library, Code 0142
Naval Postgraduate School
Monterey, CA 93943-5002 | 2 |
| 3. | Lt. Michael J. DeHaemer, Jr, USN
217 E. Montgomery St.
Baltimore, MD 21230 | 2 |
| 4. | Dr. Michael J. Zyda
Naval Postgraduate School
Code 52, Department of Computer Science
Monterey, CA 93943-5100 | 20 |
| 5. | Man-Tak Shing
Naval Postgraduate School
Code 52, Department of Computer Science
Monterey, CA 93943-5100 | 1 |
| 6. | John Maynard
Naval Ocean Systems Center
Code 402
San Diego, CA 92152 | 1 |
| 7. | Duane Gomez
Naval Ocean Systems Center
Code 433
San Diego, CA 92152 | 1 |

8. James R. Louder 1
Naval Underwater Systems Center
Combat Control Systems Department
Building 1171/1,
Newport, RI 02841

6016-585

Thesis

D2463 DeHaemer

c.1 Simplification of objects rendered by polygonal approximation.

Thesis

D2463 DeHaemer

c.1 Simplification of objects rendered by polygonal approximation.

thes2463

Simplification of objects rendered by po



3 2768 000 91520 1

DUDLEY KNOX LIBRARY